

REMARKS

STATUS OF CLAIMS

Claims 1-7 and 15-21 remain in this application. Claims 22-28 have been added. Claims 8-14 have been canceled. Claims 1 and 5 have been amended.

CLAIM OBJECTIONS

As a preliminary matter, it is noted that the Office Action Summary erroneously indicated that all of Claims 1-21 were rejected, when, in fact, paragraph 9 of the Office Action merely objected to Claims 5 and 12 as being dependent upon a rejected base claim. Claim 5 has been rewritten in independent form including all the limitations of the claim from which it formerly depended and the limitations of any intervening claims. Claim 5 therefore is patentable.

Claim 5 has been amended solely for the purpose of placing Claim 5 in independent form. Therefore, Claim 5 does not include any limitations beyond those limitations that Claim 5 included before being amended.

CLAIM REJECTIONS – 35 U.S.C. § 103

The Office Action rejected Claims 1 and 8 under 35 U.S.C. § 103(a) as being unpatentable over Dasgupta, U.S. Patent No. 5,612,865 (“Dasgupta”). The Office Action rejected Claims 2, 4, 6, 9, 11, 13, 15, 17-19, and 21 under 35 U.S.C. § 103(a) as being unpatentable over Dasgupta in view of Shrivastava et al., U.S. Patent No. 6,449,734 (“Shrivastava”). The Office Action rejected Claims 3 and 10 under 35 U.S.C. § 103(a) as being unpatentable over Dasgupta in view of Wolff, U.S. Patent No. 6,185,601 (“Wolff”). The Office Action rejected Claim 16 under 35 U.S.C. § 103(a) as being unpatentable over Dasgupta in view

of Shrivastava and Wolff. The Office Action rejected Claims 7 and 14 under 35 U.S.C. § 103(a) as being unpatentable over Dasgupta in view of Freitas et al., U.S. Patent No. 6,401,110 (“Freitas”). Claims 8-14 have been canceled, thereby obviating the rejections of those claims. With respect to Claims 1-7 and 15-21, the rejections are traversed, respectfully, in view of the present amendments and remarks.

CLAIM 1

Claim 1 recites a method for re-distributing, over a cluster of one or more active nodes, management of locks on shared resources. The method includes:

establishing a first master node as master for one or more resources in response to a hash value range being mapped to said first master node, wherein the hash value range is associated with said one or more resources by a hash function;
transferring responsibility for mastering said one or more resources from the first master node to a second master node during a transfer time interval;
during the transfer time interval, a receiving node receiving new lock requests, wherein said receiving node is one of the first master node and the second master node; and
during the transfer time interval, processing the new lock requests

(Emphasis added).

To ensure that locks are properly granted, conventional systems (1) suspend all new lock operations, (2) “quiesce” by processing all outstanding lock operations, and (3) after all outstanding requests have been processed, begin the ownership transfer **while all new lock operations remain suspended**. Thus, no new lock requests are received or processed during the transfer time interval. The suspension of lock operations temporarily halts some system functions and adversely affects system performance. The suspension of lock operations is called “freezing” lock requests. However, because the receiving node of the method of Claim 1 **receives new lock requests during the transfer time interval**, the method of Claim 1 avoids

the freezing of lock requests. Thus, the method of Claim 1 avoids the adverse effects associated with freezing lock requests.

Dasgupta does not teach or suggest the method of Claim 1. Instead, Dasgupta discloses a system in which nodes do NOT receive new lock requests during the transfer time interval. Dasgupta refers to the sequence of events that occur within a distributed lock manager (“DLM”) as a result of a change of cluster membership as the “cluster reconfiguration sequence” (col. 3, lines 53-57). These events include the redistribution of the management of locks to functional nodes of a cluster (col. 3, lines 46-47). Dasgupta discloses that **before** cluster reconfiguration can proceed (col. 7, lines 37-38), all client processes are **stopped from posting any additional requests** to the DLM (col. 7, lines 53-54), and are **paused for the remainder of the cluster reconfiguration activity** (col. 7, lines 46-47). Thus, during the cluster reconfiguration activity, in which the management of locks is redistributed among nodes, NO node receives any new lock requests. In other words, the approach disclosed by Dasgupta (the “Dasgupta approach”) freezes lock requests during the transfer time interval. As a result, the Dasgupta approach fails to avoid the adverse effects avoided by the method of Claim 1.

Thus, Dasgupta fails to teach or suggest “**during the transfer time interval, a receiving node receiving new lock requests**, wherein said receiving node is one of the first master node and the second master node” as required by the method of Claim 1. In fact, because Dasgupta **teaches away** from any node receiving new lock requests during the transfer time interval, Dasgupta cannot possibly be said to inherently disclose “a receiving node receiving new lock requests during the transfer time interval.” For at least the same reasons, the method of Claim 1 would not have been obvious to a person of ordinary skill in the art, even if that person

knew what Dasgupta disclosed. For at least the reasons stated above, Applicant respectfully requests the withdrawal of the rejection of Claim 1.

CLAIMS 2-4, 6, AND 7

Claims 2-4, 6, and 7 depend from Claim 1 and therefore also include the limitation **“during the transfer time interval, a receiving node receiving new lock requests, wherein said receiving node is one of the first master node and the second master node.”** Although the Office Action rejects various ones of Claims 2-4, 6, and 7 based on various combinations of Dasgupta with various other references (i.e., Shrivastava, Wolff, and Freitas), the Office Action does not even allege that any of these various other references discloses this limitation. As discussed above, Dasgupta itself fails to teach or suggest this limitation.

Therefore, even if Dasgupta could be combined with **all** of these various other references, such a combination would still fail to teach or suggest any of the methods of Claims 2-4, 6, and 7. Such a combination would still fail to teach or suggest **“during the transfer time interval, a receiving node receiving new lock requests, wherein said receiving node is one of the first master node and the second master node”** as included in each of Claims 2-4, 6, and 7. For at least the reasons stated above, Applicant respectfully requests the withdrawal of the rejection of Claims 2-4, 6, and 7.

CLAIM 15

Claim 15 recites a computer-readable medium bearing instructions for re-distributing, over a cluster of one or more active nodes, management of locks on shared resources from a

first master node to a second master node during a transfer time interval. The instructions are arranged to cause one or more processors on the first master node to perform:

re-mapping a hash value range initially assigned to the first master node to the second master node, wherein the hash value range is associated with one or more of the shared resources by a hash function;
sending initial lock information resident on the first master node at a start of the transfer time interval to the second master node;
receiving acknowledgments at the first master node from other active nodes in the cluster; said acknowledgements indicating that said other active nodes have been informed that said second master node is assuming responsibility for mastering said one or more resources

(Emphasis added).

Dasgupta does not teach or suggest the computer-readable medium of Claim 15. In fact, the Office Action admits that Dasgupta does not teach or suggest the limitation "**receiving acknowledgments at the first master node from other active nodes in the cluster; said acknowledgements indicating that said other active nodes have been informed that said second master node is assuming responsibility for mastering said one or more resources.**"

Shrivastava also does not teach or suggest the computer-readable medium of Claim 15. The Office Action alleges that Shrivastava discloses the above limitation by disclosing that systems in a cluster have the same view of cluster membership, and in the event that one system detects a communication failure with another system, the detecting system broadcasts a message to the cluster causing other members to verify their view of the current cluster membership (col. 5, lines 46-51), and by disclosing that each node knows which nodes own which resources and groups (col. 8, lines 62-63).

However, Claim 15 requires that the instructions borne on the computer-readable medium be arranged to cause **the first master node** to receive acknowledgements indicating that other active nodes have been informed that the second master node is assuming responsibility for one

or more resources. The first master node is the node from which the management of locks is being redistributed. In contrast, the broadcast message disclosed in Shrivastava appears to be designed to cause members of a cluster to verify their views of the cluster membership. There is no teaching or suggestion in Shrivastava that verifying a view of cluster membership necessarily involves sending an acknowledgement of that verification to any particular member. In verifying its view of cluster membership, a member most likely simply checks to make sure that its view is accurate. There is no implication in Shrivastava that a cluster member verifies its view “to” another member. Thus, there is no teaching or suggestion in Shrivastava of any sort of acknowledgement being sent to a specific node from which the management of locks is being redistributed.

Furthermore, even if each of Shrivastava’s nodes knows which nodes own which resources and groups, it does not necessarily follow that any sort of acknowledgements were involved in the process by which those nodes obtained that knowledge. Shrivastava discloses a broadcast message, but Shrivastava does not disclose any acknowledgements; especially not the specific acknowledgments required by Claim 15.

Thus, Shrivastava also fails to teach or suggest “receiving acknowledgments at the first master node from other active nodes in the cluster; said acknowledgements indicating that said other active nodes have been informed that said second master node is assuming responsibility for mastering said one or more resources” as required by Claim 15.

Thus, even if Dasgupta and Shrivastava could be combined, such a combination would still fail to teach or suggest the computer-readable medium of Claim 15. Such a combination would still fail to teach or suggest **“receiving acknowledgments at the first master node**

from other active nodes in the cluster; said acknowledgements indicating that said other active nodes have been informed that said second master node is assuming responsibility for mastering said one or more resources” as required by Claim 15. For at least the reasons stated above, Applicant respectfully requests the withdrawal of the rejection of Claim 15.

CLAIM 18

Claim 18 recites a computer-readable medium bearing instructions for re-distributing, over a cluster of one or more active nodes, management of locks on shared resources from a first master node to a second master node during a transfer time interval. The instructions are arranged to cause one or more processors on the second master node to perform:

receiving initial lock information resident on the first master node at a start of the transfer time interval;
re-mapping a hash value range initially assigned to the first master node to the second master node, wherein the hash value range is associated with one or more of the shared resources by a hash function; and
sending a broadcast message to all other nodes in the cluster that the second master node is a new master node for resources associated with the hash value range

(Emphasis added).

Dasgupta does not teach or suggest the computer-readable medium of Claim 18. In fact, the Office Action admits that Dasgupta does not teach or suggest the limitation **“sending a broadcast message to all other nodes in the cluster that the second master node is a new master node for resources associated with the hash value range.”**

Shrivastava also does not teach or suggest the computer-readable medium of Claim 18. The Office Action alleges that Shrivastava discloses the above limitation by disclosing that systems in a cluster have the same view of cluster membership, and in the event that one system detects a communication failure with another system, the detecting system broadcasts a message

to the cluster causing other members to verify their view of the current cluster membership (col. 5, lines 46-51), and by disclosing that each node knows which nodes own which resources and groups (col. 8, lines 62-63).

However, the broadcast message disclosed in Shrivastava appears to be designed to cause members of a cluster to verify their views of the **cluster membership**. Cluster membership is not the same as resource ownership. There is no teaching or suggestion in Shrivastava that the message that causes members to verify their views of the cluster membership also informs the members that a particular node is a new master node for resources associated with a particular hash value range. Thus, there is no teaching or suggestion in Shrivastava of any sort of broadcast message that indicates that a particular node is a new master node for resources associated with a particular hash value range.

Furthermore, even if each of Shrivastava's nodes knows which nodes own which resources and groups, it does not necessarily follow that those nodes obtain that knowledge from a broadcast message that indicates that a particular node is a new master node for resources associated with a particular hash value range. Shrivastava discloses a broadcast message, but Shrivastava does not disclose that the broadcast message indicates the specific information required by Claim 18.

Thus, Shrivastava also fails to teach or suggest "**sending a broadcast message to all other nodes in the cluster that the second master node is a new master node for resources associated with the hash value range**" as required by Claim 18.

Thus, even if Dasgupta and Shrivastava could be combined, such a combination would still fail to teach or suggest the computer-readable medium of Claim 18. Such a combination would still fail to teach or suggest "**sending a broadcast message to all other nodes in the**

cluster that the second master node is a new master node for resources associated with the hash value range” as required by Claim 18. For at least the reasons stated above, Applicant respectfully requests the withdrawal of the rejection of Claim 18.

CLAIM 21

Claim 21 recites a computer-readable medium bearing instructions for re-distributing, over a cluster of one or more active nodes, management of locks on shared resources from a first master node to a second master node during a transfer time interval. The instructions are arranged to cause one or more processors on a third node to perform:

receiving a broadcast message indicating that the second master node is a new master node for resources associated with a hash value range, wherein the hash value range is associated with one or more of the shared resources by a hash function; re-mapping the hash value range to the second master node;
sending an acknowledgment to the first master node in response to the broadcast message, said acknowledgement indicating that said third node has been informed that said second master node is assuming responsibility for mastering said one or more resources; and
after said sending an acknowledgment, sending subsequent lock requests for the one or more of the shared resources to the second master node

(Emphasis added).

Dasgupta does not teach or suggest the computer-readable medium of Claim 21. In fact, the Office Action admits that Dasgupta does not teach or suggest the limitation “**sending an acknowledgement to the first master node in response to the broadcast message, said acknowledgement indicating that said third node has been informed that said second master node is assuming responsibility for mastering said one or more resources.**”

Shrivastava also does not teach or suggest the computer-readable medium of Claim 21. The Office Action alleges that Shrivastava discloses the above limitation by disclosing that

systems in a cluster have the same view of cluster membership, and in the event that one system detects a communication failure with another system, the detecting system broadcasts a message to the cluster causing other members to verify their view of the current cluster membership (col. 5, lines 46-51), and by disclosing that each node knows which nodes own which resources and groups (col. 8, lines 62-63).

However, as is discussed above with regard to Claim 15, there is no teaching or suggestion in Shrivastava of any sort of acknowledgement being sent to a specific node from which the management of locks is being redistributed. Thus, Shrivastava also fails to teach or suggest **“sending an acknowledgment to the first master node in response to the broadcast message, said acknowledgement indicating that said third node has been informed that said second master node is assuming responsibility for mastering said one or more resources”** as required by Claim 21.

Thus, even if Dasgupta and Shrivastava could be combined, such a combination would still fail to teach or suggest the computer-readable medium of Claim 21. Such a combination would still fail to teach or suggest **“sending an acknowledgment to the first master node in response to the broadcast message, said acknowledgement indicating that said third node has been informed that said second master node is assuming responsibility for mastering said one or more resources”** as required by Claim 21. For at least the reasons stated above, Applicant respectfully requests the withdrawal of the rejection of Claim 21.

MISCELLANEOUS

The pending claims not discussed so far are dependent claims that depend on an independent claim that is discussed above. Because each of the dependent claims includes the

limitations of claims upon which they depend, the dependent claims are patentable for at least those reasons the claims upon which the dependent claims depend are patentable. Removal of the rejections with respect to the dependent claims and allowance of the dependent claims is respectfully requested. In addition, the dependent claims introduce additional limitations that independently render them patentable. Due to the fundamental difference already identified, a separate discussion of those limitations is not included at this time.

All issues raised in the Office Action have been addressed. Allowance of the pending claims is appropriate. Entry of the amendments herein is requested, respectfully.

The Examiner is invited to telephone the undersigned at (408) 414-1080 to discuss any issue that may advance prosecution.

No fee is believed to be due specifically in connection with this Reply. To the extent necessary, Applicants petition for an extension of time under 37 C.F.R. § 1.136. The Commissioner is authorized to charge any fee that may be due in connection with this Reply to our Deposit Account No. 50-1302.

I hereby certify this correspondence is being deposited with the United States Postal Service as first class mail in an envelope addressed to: Mail Stop Non-Fee Amendment, Commissioner for Patents, Box 1450, Alexandria, VA 22313-1450.

Date of deposit: October 14, 2003

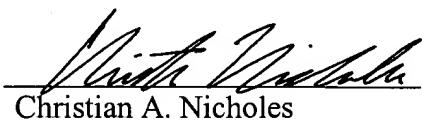
Typed Name: Jennifer Newell

Signature: Jennifer Newell

Respectfully submitted,

Hickman Palermo Truong & Becker, LLP

By


Christian A. Nicholes

Attorney for Applicant(s)

Reg. No. 50,266

Date: October 14, 2003

Address: 1600 Willow Street

San Jose, California 95125-5106

Telephone No.: (408) 414-1080